

Inferring Labor Income Risk From Economic Choices: An Indirect Inference Approach*

Fatih Guvenen[†]

Anthony Smith[‡]

Preliminary and Incomplete. Comments Welcome.

April 27, 2008

Abstract

This paper sheds light on the nature of labor income risk by exploiting the information contained in the joint dynamics of households' labor earnings and consumption-choice decisions. In particular, this paper attempts to discriminate between two leading views on the nature of labor income risk: the “restricted income profiles” (RIP) model—in which individuals are subjected to large and persistent income shocks but face similar life-cycle income profiles—and the “heterogeneous income profiles” (HIP) model—in which individuals are subjected to income shocks with modest persistence but face individual-specific income profiles. Although these two different income processes have vastly different implications for economic behavior, earlier studies have found that labor income data alone is insufficient to distinguish between them. This paper, therefore, brings to bear the information embedded in consumption data. Specifically, we apply the powerful new tools of indirect inference to rich panel data on consumption and labor earnings to estimate a rich structural consumption-savings model. The method we develop is very flexible and allows the estimation of income processes from economic decisions in the presence of non-separabilities between consumption and leisure, partial insurance of income shocks, frequently binding borrowing constraints, missing observations, among others. In this estimation, we use an auxiliary model—which forms the bridge between the data and the consumption-savings model—that provides a sharp distinction between the RIP and HIP models. Finally, we conduct formal statistical tests to assess the extent to which the RIP and HIP models find support in the data.

*For helpful comments and discussions, we thank Daron Acemoglu, Victor Aguirregabiria, Orazio Attanasio, Richard Blundell, Martin Browning, Raj Chetty, Lars Hansen, Yuichi Kitamura, Luigi Pistaferri, Sam Schulhofer-Wohl, Kjetil Storesletten, Harald Uhlig, seminar participants at Georgetown, HEC Lausanne, Princeton, Seoul National, Sogang University, Stanford, Tokyo University, UCL, University of Iowa, University of Toronto, and the conference participants at the NBER Summer Institute, Banque de France/Bundesbank Spring 2007 Conference, 2007 SED Conference, 2007 Paris Workshop on Heterogeneity and Macrodynamics, IFS/UCL Workshop on Consumption Inequality and Income Dynamics, Advances in Theory Based Estimation Conference in Honor of Hansen and Singleton, CEPR's 2007 ESSIM Conference, and 2008 AEA Winter meetings. Matthew Johnson provided excellent research assistance. Guvenen acknowledges financial support from the NSF under grant number SES-0649437. All remaining errors are our own.

[†]University of Minnesota and NBER; guvenen@umn.edu; <http://www.econ.umn.edu/~guvenen>

[‡]Yale University; tony.smith@yale.edu; <http://www.econ.yale.edu/smith>.

1 Introduction

The goal of this paper is to elicit information about the nature of labor income risk from individuals' economic decisions (such as consumption-savings choice), which contain valuable information about the environment faced by individuals, including the future (income) risks they perceive.

To provide a framework for this discussion, consider the following process for log labor income of individual i with t years of labor market experience:

$$y_t^i = \underbrace{[a_0 + a_1t + a_2t^2 + a_3Educ + \dots]}_{\text{common life-cycle component}} + \underbrace{[\alpha^i + \beta^i t]}_{\text{profile heterogeneity}} + \underbrace{[z_t^i + \varepsilon_t^i]}_{\text{stochastic component}} \quad (1)$$

$$\text{where} \quad z_t^i = \rho z_{t-1}^i + \eta_t^i, \quad \text{and} \quad \eta_t^i, \varepsilon_t^i \sim iid$$

The terms in the first bracket capture the life-cycle variation in labor income that is *common* to all individuals with given observable characteristics. The second component captures potential *individual-specific* differences in income growth rates (as well as in levels, which is less important). Such differences would be implied for example by a human capital model with heterogeneity in learning ability.¹ Finally, the terms in the last bracket represent the stochastic variation in income, which is written here as the sum of an AR(1) component and a purely transitory shock. This is a specification commonly used in the literature.

A vast empirical literature has estimated various versions of (1) in an attempt to answer the following two questions:

1. Do individuals differ *systematically* in their income growth rates? If such differences exist, are they quantitatively important? i.e., is $\sigma_\beta^2 \gg 0$?
2. How large and how persistent are income shocks? i.e., what is σ_η^2 and ρ ?

Existing studies in the literature can be broadly categorized into two groups based on the conclusions they reach regarding these questions. The first group of papers *impose* $\sigma_\beta^2 \equiv 0$ based on outside evidence,² and with this restriction estimate ρ to be close to 1. We refer to this version of the process in (1) as the “*Restricted Income Profiles*” (*RIP*) model. The second group of papers do not impose any restrictions on (1) and find that ρ is significantly less than 1 and σ_β^2 is large. We

¹See for example, the classic paper by Ben-Porath (1967). For more recent examples of such a human capital model, see Guvenen and Kuruscu (2006), and Huggett, Ventura and Yaron (2006).

²The outside evidence refers to a test proposed by MaCurdy (1982) in which he failed to reject the null of RIP against the alternative of HIP. Two recent papers, Baker (1997) and Guvenen (2007b), argue that tests based on average autocovariances lack power against the alternative of a HIP process with an autoregressive component, and therefore, the lack of rejection of the RIP null does not provide evidence against the HIP model.

refer to this version of (1) as the “*Heterogeneous Income Profiles*” (*HIP*) model. In other words, according to the RIP view, most of the rise in within-cohort income inequality over the life-cycle is due to large and persistent shocks, whereas in the HIP view, it is due to systematic differences in income growth rates. While overall we interpret the results of these studies, and especially those of the more recent papers, as more supportive of the HIP model, it is fair to say that this literature has not produced an unequivocal verdict.³

A key point to observe is that these existing studies do not utilize the information revealed by individuals’ consumption-savings choice to distinguish between the HIP and RIP models.⁴ But endogenous choices, such as consumption and savings, contain valuable information about the environment faced by individuals, including the future risks they perceive. Therefore, the main purpose of this paper is to use the restrictions imposed by the RIP and HIP processes on consumption data—in the context of a life-cycle model—to bring more evidence to bear on this important question. We elaborate further below on the advantages of focusing on consumption-savings choice (instead of using labor income data in isolation or using other endogenous choices, such as labor supply) for drawing inference about the labor income process.

In a sense, the two questions discussed so far only scratch the surface of “the nature of income risk.” This is because those two questions are statistical in nature, i.e., they relate to how the income process is viewed by the *econometrician* who studies past observations on individual income. But it is quite plausible that individuals may have more, or less, information about their income process than the econometrician at different points in their lifecycle, which raises two more questions:

3. If individuals indeed differ in their income growth rate as suggested by the HIP model, how much do individuals *know* about their β^i at *different points* in their life-cycle? In other words, what fraction of the heterogeneity in β^i constitutes “uncertainty” on the part of individuals as opposed to simply being some “known heterogeneity”?
4. What fraction of income movements measured by z_t^i and ε_t^i are really “unexpected shocks” as opposed to being “anticipated changes”?

³A short list of these studies includes MaCurdy (1982), Abowd and Card (1989), and Topel (1990), which find support for the RIP model; Lillard and Weiss (1979), Hause (1980), and especially the more recent studies such as Baker (1997), Haider (2001), and Guvenen (2007b) which find support for the HIP model.

⁴Two recent papers do use consumption data but in a more limited fashion than this paper intends to do. In a recent paper, Huggett, Ventura and Yaron (2006) study a version of the Ben-Porath model and make some use of consumption data to measure the relative importance of persistent income shocks versus heterogeneity in learning ability. Although the income process generated by their model does not exactly fit into the specification in equation (1) their results are informative. Second, Guvenen (2007a) uses consumption data to investigate if a HIP model estimated from income data is consistent with some stylized consumption facts. While both of these papers are informative about the HIP versus RIP debate, they make limited use of consumption data, especially of the *dynamics* of consumption behavior.

These questions are inherently different than the first two in that they pertain to how *individuals* perceive their income process. As such, they cannot be answered using income data alone, but the answers can be teased out, again, from individuals' economic decisions. To give one example (to question 4), consider a married couple who jointly decide that they will both work up to a certain age and then will have children at which time one of the spouses will quit his/her job to take care of the children. The ensuing large fall in household income will appear as a large permanent shock to the econometrician using labor income data alone, but consumption (and savings) data would reveal that this change has been anticipated.

Several papers have used consumption data and shed light on various properties of income processes (among others, Hall and Mishkin (1982), Deaton and Paxson (1994), Blundell and Preston (1998), and Blundell, Pistaferri and Preston (2006a)). This paper aims to contribute to this literature in the following ways. First and foremost, existing studies consider only versions of the RIP model (i.e., they set $\sigma_\beta^2 \equiv 0$ at the outset), whereas our goal is to distinguish between HIP and RIP models as well. Second, and furthermore, these studies also *impose* $\rho \equiv 1$, and only estimate the innovation variances. In other words, there is no existing study to our knowledge that uses consumption data and estimates ρ . Therefore, this paper will leave ρ unrestricted (even in the RIP version) and exploit consumption and income data jointly to pin down its value. Since many incomplete markets models are still calibrated using versions of the RIP process, the results of this exercise should be useful for calibrating those models. The third contribution of this paper will be in the method used for estimation—indirect inference—which is much less restrictive than, and has several important advantages over, the GMM approach used in previous work.⁵

1.1 Why Look at Consumption-Savings Choice?

Even if one is only interested in the first two questions raised above, using information revealed by intertemporal choices has important advantages. This is because one difficulty of using income data alone is that identification between HIP and RIP models partly depends on the behavior of the *higher-order* autocovariances of income.

To see this clearly, consider the case where the panel data set contains income observations on a single cohort over time. In this case, the second moments of the cross-sectional distribution for this cohort are given by:

$$\begin{aligned} \text{var}(y_t^i) &= [\sigma_\alpha^2 + 2\sigma_{\alpha\beta}t + \sigma_\beta^2 t^2] + \text{var}(z_t^i) + \sigma_\varepsilon^2 \\ \text{cov}(y_t^i, y_{t+n}^i) &= [\sigma_\alpha^2 + \sigma_{\alpha\beta}(2t+n) + \sigma_\beta^2 t(t+n)] + \rho^n \text{var}(z_t^i), \end{aligned} \tag{2}$$

⁵Two important differences of the present paper from Guvenen (2007a) is that that paper (i) only estimated $\hat{\sigma}_{\beta|0}^2$ from consumption data, taking all other parameters as estimated from income data, and (ii) only used the rise in within-cohort consumption inequality as a moment condition. The present paper instead (i) brings consumption data to bear on the estimation of the *entire* vector of structural parameters, and (ii) does this by systematically focusing on the dynamic relationship between consumption and income movements.

where $t = 1, \dots, T$, and $n = 1, \dots, T - t$. There are two sources of identification between the RIP and HIP processes, which can be seen by inspecting these formulas. The first piece of information is provided by the change in the cross-sectional variance of income as the cohort ages (i.e., the diagonal elements of the variance-covariance matrix), which is shown on the first line of (2). The terms in the square bracket capture the effect of profile heterogeneity, which is a *convex* increasing function of age. The second term captures the effect of the AR(1) shock, which is a *concave* increasing function of age as long as $\rho < 1$. Thus, if the variance of income in the data increases in a convex fashion as the cohort gets older, this would be captured by the HIP terms (notice that the coefficient on t^2 is σ_β^2), whereas a non-convex shape would be captured by the presence of AR(1) shocks.

The second source of identification is provided by the autocovariances displayed in the second line. The covariance between ages t and $t + n$ is again composed of two parts. As before, the terms in the square bracket capture the effect of heterogeneous profiles and is a convex function of age. Moreover, the coefficients of the linear and quadratic terms depend both on t and n , which allows covariances to be decreasing, increasing or non-monotonic in n at each t . The second term captures the effect of the AR(1) shock, and notice that for a given t , it depends on the covariance lag n only through the geometric discounting term ρ^n . The strong prediction of this form is that, starting at age t , covariances should decay geometrically at the rate ρ , regardless of the initial age. Thus, in the RIP model (which only has the AR(1) component) covariances are restricted to decay at the same rate at every age, and cannot be non-monotonic in n .

Notice that for a cohort with 40 years of working life, there are only 40 variance terms, but many more—780 ($= (40 \times 41) / 2 - 40$) to be precise—autocovariances, which provide crucial information for distinguishing between HIP and RIP processes. The main difficulty is that because of sample attrition, fewer and fewer individuals contribute to these higher autocovariances, raising important concerns about potential selectivity bias. To give a rough idea, if one uses labor income data from the Panel Study of Income Dynamics (PSID), and selects all individuals who are observed in the sample for 3 years or more (which is a typical sample selection criterion), the number of individuals contributing to the 20th autocovariance will be about 1/5 of the number of individuals contributing to the 3rd autocovariance. To the extent that these individuals are not a completely random subsample of the original sample, covariances at different lags will have variation due to sample selection that can confound the identification between HIP and RIP models.

In contrast, because of its forward-looking nature, even short-run movements in consumption, and the immediate response of consumption to income innovations contain information about the perceived *long-run* behavior of the income process. Therefore even lower-order covariances of consumption would help in distinguishing HIP from RIP. (Notice that the dynamic aspect of the consumption-savings choice also distinguishes it from other decisions, such as labor supply, which are static in nature, unless one models intertemporally non-separable preferences in leisure.)

2 Bayesian Learning about Income Profiles

Embedding the HIP process into a life-cycle model requires one take a stand on what individuals know about their own β^i . We follow Guvenen (2007a) and assume that individuals enter the labor market with some prior belief about their β^i and then update their beliefs over time in a Bayesian fashion. Notice that the prior variance of this belief (denote by $\hat{\sigma}_{\beta|0}^2$) measures how uncertain individuals are about their own β^i at time zero, addressing question 3 above.

We now cast the learning process as a Kalman filtering problem which allows us to obtain recursive updating formulas for beliefs. Individuals (know α^i), observe y_t^i , and must learn about $\mathbf{S}_t^i \equiv (\beta^i, z_t^i)$.⁶ It is convenient to express the learning process as a Kalman filtering problem using the state-space representation. In this framework, the “state equation” describes the evolution of the vector of state variables that is unobserved by the decision maker:

$$\underbrace{\begin{bmatrix} \beta^i \\ z_{t+1}^i \end{bmatrix}}_{\mathbf{S}_{t+1}^i} = \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & \rho \end{bmatrix}}_{\mathbf{F}} \underbrace{\begin{bmatrix} \beta^i \\ z_t^i \end{bmatrix}}_{\mathbf{S}_t^i} + \underbrace{\begin{bmatrix} 0 \\ \eta_{t+1}^i \end{bmatrix}}_{\boldsymbol{\nu}_{t+1}^i}$$

Even though the parameters of the income profile have no dynamics, including them into the state vector yields recursive updating formulas for beliefs using the Kalman filter. A second (observation) equation expresses the observable variable(s) in the model—in this case, log income—as a linear function of the underlying hidden state and a transitory shock:

$$y_t^i = \alpha^i + \begin{bmatrix} t & 1 \end{bmatrix} \begin{bmatrix} \beta^i \\ z_t^i \end{bmatrix} + \varepsilon_t^i = \alpha^i + \mathbf{H}_t' \mathbf{S}_t^i + \varepsilon_t^i$$

We assume that both shocks have *i.i.d* Normal distributions and are independent of each other, with \mathbf{Q} and R denoting the covariance matrix of $\boldsymbol{\nu}_t^i$ and the variance of ε_t^i respectively. To capture an individual’s initial uncertainty, we model his prior belief over (β^i, z_1^i) by a multivariate Normal distribution with mean

$$\hat{\mathbf{S}}_{1|0}^i \equiv (\hat{\beta}_{1|0}^i, \hat{z}_{1|0}^i)$$

and variance-covariance matrix:

$$\mathbf{P}_{1|0} = \begin{bmatrix} \sigma_{\beta,0}^2 & 0 \\ 0 & \sigma_{z,0}^2 \end{bmatrix}$$

where we use the short-hand notation $\sigma_{\cdot,t}^2$ to denote $\sigma_{\cdot,t+1|t}^2$. After observing $(y_t^i, y_{t-1}^i, \dots, y_1^i)$, the posterior belief about \mathbf{S}^i is Normally distributed with a mean vector $\hat{\mathbf{S}}_t^i$, and covariance matrix \mathbf{P}_t . Similarly, let $\hat{\mathbf{S}}_{t+1|t}^i$ and $\mathbf{P}_{t+1|t}$ denote the one-period-ahead *forecasts* of these two variables

⁶Guvenen (2007a) also allows for learning about α^i and shows that it has a minimal effect on the behavior of the model. Therefore, we abstract from this feature which eliminates one state variable and simplifies the problem.

respectively. These two variables play central roles in the rest of our analysis. Finally, log income has a Normal distribution conditional on an individual's beliefs:

$$y_{t+1}^i | \widehat{\mathbf{S}}_t^i \sim N \left(\mathbf{H}'_{t+1} \widehat{\mathbf{S}}_{t+1|t}^i, \mathbf{H}'_{t+1} \mathbf{P}_{t+1|t} \mathbf{H}_{t+1} + R \right). \quad (3)$$

In this particular problem, the standard Kalman filtering equations can be manipulated to obtain some simple expressions that will become useful later. To this end, define:

$$\begin{aligned} A_t &\equiv t\sigma_{\beta,t|t-1}^2 + \sigma_{\beta z,t|t-1}, \\ B_t &\equiv t\sigma_{\beta z,t|t-1} + \sigma_{z,t|t-1}^2, \\ X_t &\equiv \text{var}_{t-1}(y_t^i) = A_t t + B_t + R \end{aligned}$$

Using the Kalman recursion formulas:

$$\begin{bmatrix} \widehat{\beta}_{t+1|t}^i \\ \widehat{z}_{t+1|t}^i \end{bmatrix} = \begin{bmatrix} \widehat{\beta}_{t|t-1}^i \\ \rho \widehat{z}_{t|t-1}^i \end{bmatrix} + \begin{bmatrix} A_t/X_t \\ \rho B_t/X_t \end{bmatrix} \left(y_t^i - \left(\widehat{\beta}_{t|t-1}^i t + \widehat{z}_{t|t-1}^i \right) \right)$$

Define the innovation to beliefs:

$$\widehat{\xi}_t = y_t^i - \left(\widehat{\beta}_{t|t-1}^i t + \widehat{z}_{t|t-1}^i \right)$$

Then we can rewrite:

$$\widehat{\beta}_{t+1|t}^i - \widehat{\beta}_{t|t-1}^i = (A_t/X_t) \widehat{\xi}_t \quad (4)$$

$$\widehat{z}_{t+1|t}^i - \rho \widehat{z}_{t|t-1}^i = (\rho B_t/X_t) \widehat{\xi}_t \quad (5)$$

An important point to note is that $\widehat{\xi}_t$ and (the true innovation to income) η_t^i do not need to have the same sign, a point that will play a crucial role below. Finally, the posterior variances evolve:

$$\sigma_{\beta,t+1|t}^2 = \sigma_{\beta,t|t-1}^2 - \frac{A_t^2}{X_t} \quad (6)$$

$$\sigma_{z,t+1|t}^2 = \rho^2 \left[\sigma_{z,t|t-1}^2 - \frac{B_t^2}{X_t} \right] + R \quad (7)$$

For a range of parameterizations A/X has an inverse U-shape over the life-cycle. Therefore, beliefs about β^i changes (and precision rises) slowly early on but become faster over time. In contrast, B/X declines monotonically. As shown in Guvenen (2007a), optimal learning in this model has some interesting features. In particular, learning is very slow and the speed of learning has a non-monotonic pattern over the life-cycle (which is due to the fact that A/X has an inverse

U-shape). If instead the prior uncertainty were to resolve quickly, consumption behavior after the first few years would not be informative about the prior uncertainty faced by individuals ($\widehat{\sigma}_{\beta|0}^2$).

Finally we discuss how an individual's prior belief about β^i is determined. Suppose that the distribution of income growth rates in the population is generated as $\beta^i = \beta_k^i + \beta_u^i$, where β_k^i and β_u^i are two random variables, independent of each other, with zero mean and variances of $\sigma_{\beta_k}^2$ and $\sigma_{\beta_u}^2$. Clearly then, $\sigma_{\beta}^2 = \sigma_{\beta_k}^2 + \sigma_{\beta_u}^2$. The key assumption we make is that individual i observes the realization of β_k^i , but not of β_u^i (hence the subscripts indicate *known* and *unknown*, respectively). Under this assumption, the prior mean of individual i is $\widehat{\beta}_{1|0}^i = \beta_k^i$, and the prior variance is $\sigma_{\beta,0}^2 = \sigma_{\beta_u}^2 = (1 - \lambda)\sigma_{\beta}^2$, where we define $\lambda = 1 - \sigma_{\beta_u}^2/\sigma_{\beta}^2$, as the fraction of variance known by individuals. Two polar cases deserve special attention. If $\lambda = 0$, individuals do not have any *private prior information* about their income growth rate (i.e., $\sigma_{\beta,0}^2 = \sigma_{\beta}^2$ and $\widehat{\beta}_{1|0}^i = \bar{\beta}$ for all i , where $\bar{\beta}$ is the population average). On the other hand if $\lambda = 1$, each individual observes β^i completely and faces no prior uncertainty about its value.

2.1 The HIP Model

Consider an environment where each individual lives for T years and works for the first R ($< T$) years of his life, after which he retires. Individuals do not derive utility from leisure and hence supply labor inelastically.⁷ During the working life, the income process is given by the HIP process specified in equation (1). During retirement, the individual receives a pension which is given by a fixed fraction of the individual's income in period R .⁸ There is a risk-free bond that sells at price P^b (with a corresponding net interest rate $r^f \equiv 1/P^b - 1$). Individuals can also borrow at the same interest rate up to an age-specific borrowing constraint \underline{W}_{t+1} , specified below.

The relevant state variables for this dynamic problem are the asset level, ω_t^i , and his current forecast of the true state in the current period, $\widehat{\mathbf{S}}_t$. The dynamic programming problem of the individual can be written as:

$$V_t^i(\omega_t^i, \widehat{\mathbf{S}}_t^i) = \max_{C_t^i, \omega_{t+1}^i} \left\{ U(C_t^i) + \delta E_t \left[V_{t+1}^i(\omega_{t+1}^i, \widehat{\mathbf{S}}_{t+1}^i) \right] \right\}$$

s.t.

$$C_t^i + a_{t+1}^i = \omega_t \tag{8}$$

$$\omega_t = (1 + r) a_t^i + Y_t^i \tag{9}$$

$$a_{t+1}^i \geq \underline{W}_{t+1}, \quad \text{and}$$

and Kalman recursions

⁷The labor supply choices of both the husband and wife appear to be important for drawing robust inference about the nature of income risk. Therefore, we intend to introduce labor supply choice for both spouses in future versions of this paper. Such extensions are conceptually feasible with indirect inference, although it increases computational costs.

⁸A more realistic Social Security system will be introduced in a later version of the paper.

for $t = 1, \dots, R - 1$, where $Y_t^i \equiv e^{y_t^i}$ is the *level* of income, and V_t^i is the value function of a t year-old individual. The evolutions of the vector of beliefs and its covariance matrix are governed by the Kalman recursions given in equations (4, 5, 6, 7). Finally, the expectation is taken with respect to the conditional distribution of y_{t+1}^i given by equation (3), since this is the source of all uncertainty in the model.

During retirement, pension income is constant and since there is no other source of uncertainty or learning, the problem simplifies significantly:

$$\begin{aligned} V_t^i(\omega_t^i, Y^i) &= \max_{c_t^i, \omega_{t+1}^i} [U(C_t^i) + \delta V_{t+1}^i(\omega_{t+1}^i, Y^i)] \\ \text{s.t.} \quad Y^i &= \Phi(Y_R^i), \text{ and eq. (8, 9)} \end{aligned} \tag{10}$$

for $t = R, \dots, T$, and $V_{T+1} \equiv 0$.

2.2 The RIP Model

The second model is essentially the same as the first one, with the exception that the income process is now given by a RIP process. Because with a RIP process all individuals share the same life-cycle income profile (α, β) , there is no learning about individual profiles, the problem simplifies significantly. Specifically, the dynamic programming problem of a typical worker is:

$$\begin{aligned} J_t^i(\omega_t^i, z_t^i) &= \max_{c_t^i, \omega_{t+1}^i} \{U(c_t^i) + \delta E [J_{t+1}^i(\omega_{t+1}^i, z_{t+1}^i) | z_t^i]\} \\ \text{s.t.} \quad &\text{equations (8, 9)} \end{aligned}$$

for $t = 1, \dots, R - 1$, where J_t^i is the value function of a t year-old individual. Notice that we assume the worker observes the persistent component of the income process, z_t^i , separately from y_t^i . This is the standard assumption in the existing consumption literature which uses the RIP process, and we follow them for comparability. Finally, because there is no income risk after retirement, the problem of a retiree is the same as in (10) above.

Notice that the HIP model does *not* nest the RIP model described here, although it comes quite close. In particular, when $\sigma_\beta^2 \equiv 0$ the HIP process does reduce to the RIP process, but now in the consumption-savings model individuals are assumed not to observe the AR(1) shock and the i.i.d shock separately (whereas in the RIP model described here, they do). We choose the RIP model not nested in the HIP model because it corresponds more closely to the framework studied in the consumption literature.

2.3 Modeling Partial Insurance

[To be added]

2.4 Introducing Endogenous Labor Supply Choice

[To be added]

3 An Indirect Inference Approach

Indirect inference is a simulation-based method for estimating, or making inferences about, the parameters of economic models. It is most useful in estimating models for which the likelihood function (or any other criterion function that might form the basis of estimation) is analytically intractable or too difficult to evaluate, as is the case here: neither one of the consumption-savings models described above yields simple estimable equations that would allow a maximum likelihood or GMM estimation. Previous studies (which focused only on the RIP model) made a number of simplifying assumptions, such as the absence of binding borrowing constraints, separability between consumption and leisure in the utility function, a simplified retirement structure, and so on, and employed several approximations to the true structural equations in order to make GMM feasible.

Instead, the hallmark of indirect inference is the use of an “auxiliary model” to capture aspects of the data upon which to base the estimation. One key advantage of indirect inference over GMM is that this auxiliary model does *not* need to correspond to any valid moment condition of the structural model for the estimates of the structural parameters to be consistent. This allows significant flexibility in choosing an auxiliary model: it can be any sufficiently rich statistical model relating the model variables to each other as long as each structural parameter of the economic model has an independent effect on the likelihood of the auxiliary model.⁹ This also allows one to incorporate many realistic features into the structural model without having to worry about whether or not one can directly derive the likelihood (or moment conditions for GMM) in the presence of these features.

While indirect inference shares a basic similarity to MSM (Method of Simulated Moments), it differs from MSM in its use of an auxiliary model to form “moment conditions”. In particular, indirect inference allows one to think in terms of structural and dynamic relationships of economic models that are difficult to express as simple unconditional moments as is often done with MSM. We illustrate this in the description of the auxiliary model below.

⁹In addition to some regularity conditions that the auxiliary model has to satisfy the precise specification of the auxiliary model will also matter for the efficiency of the estimator.

3.1 Towards an Auxiliary Model

To understand the auxiliary model that will be used, it is useful to elaborate on the dependence of consumption choice on income shocks. As noted above, the key idea behind an auxiliary model is that it should be an econometric model that is easy to estimate, yet one that captures the key statistical relations between the variables of interest in the model. Good candidates for an auxiliary model are provided by structural relationships that hold in models that are similar to the HIP and RIP models described above, and yet simple enough to allow the derivation of such relationships.

To this end, consider a simplified version of the HIP model, where we assume: (i) quadratic utility; (ii) $\delta(1+r^f) = 1$, and (iii) no retirement. Further consider a simpler form of the HIP process:

$$Y_t^i = \alpha^i + \beta^i t + z_t^i, \quad (11)$$

where income, instead of its logarithm, is linear in the underlying components, and we set $\varepsilon_t^i \equiv 0$. Under these assumptions, optimal consumption choice implies

$$\Delta C_t = \frac{1}{\varphi_t} \left[(1-\gamma) \sum_{s=0}^{T-t} \gamma^s (E_t - E_{t-1}) Y_{t+s}^i \right], \quad (12)$$

where $\gamma = 1/(1+r^f)$ and $\varphi_t = (1-\gamma^{T-t+1})$ is the annuitization factor. Substituting the simple HIP process in (11), we have:

$$E_t(Y_{t+s}^i) = \alpha^i + \hat{\beta}_t^i(t+s) + \rho^s \hat{z}_t^i$$

$$(E_t - E_{t-1})Y_{t+s}^i = (\hat{\beta}_t^i - \hat{\beta}_{t-1}^i)(t+s) + \rho^s \hat{\eta}_t^i$$

Substituting this into (12), one can show:

$$\Delta C_t = \Phi_{t,T}^r (\hat{\beta}_t^i - \hat{\beta}_{t-1}^i) + \Psi_{T-t}^{r,\rho} \hat{\eta}_t^i \quad (13)$$


where:

$$\begin{aligned} \Phi_{t,T}^r &\equiv \left[\left(\frac{\gamma}{1-\gamma} \right) + \frac{t - (T+1)\gamma^{T-t+1}}{1-\gamma^{T-t+1}} \right] \\ \Psi_{T-t}^{r,\rho} &\equiv \frac{1-\gamma}{1-\gamma\rho} \left[\frac{(1-(\gamma\rho)^{T-t+1})}{(1-\gamma^{T-t+1})} \right] \end{aligned}$$

Note that $\Phi_{t,T}^r$ is a (known) slightly convex increasing function of t , and $\Psi_{T-t}^{r,\rho}$ is constant and equal to 1 when $\rho = 1$. Recall that the Kalman filtering formulas above implied:

$$\begin{aligned} \hat{\beta}_t^i - \hat{\beta}_{t-1}^i &= (A_t/X_t) \hat{\xi}_t \\ \hat{z}_t^i - \rho \hat{z}_{t-1}^i &= (B_t/X_t) \hat{\xi}_t \end{aligned} \quad (14)$$

Figure 1: Distinguishing HIP from RIP (from Consumption Changes)



which is obtained easily from equations (4), but now $\widehat{\xi}_t$ has to be reinterpreted as the level deviation: $Y_t^i - (\widehat{\beta}_{t|t-1}^i t + \widehat{z}_{t|t-1}^i)$. Plugging this, we get in the HIP model:

$$\Delta C_t = [\Phi_{t,T}^r (A_t/X_t) + \Psi_{T-t}^{r,\rho} (\rho^s B_t/X_t)] \times \widehat{\xi}_t \quad (15)$$

Instead in the RIP model we have:


$$\Delta C_t = \Psi_{T-t}^{r,\rho} \times \eta_t^i \quad (16)$$

The last two equations underscore the key difference between the two frameworks: in the RIP model only current η_t^i matters for consumption response, whereas in the HIP model the entire history of shocks matters.¹⁰ As a result, two individuals hit by the same η_t^i may react differently depending on their history. Specifically, in the HIP model η_t^i and $\widehat{\xi}_t$ may have different signs. Therefore, an increase in income ($\Delta Y_t^i > 0$) may cause a fall in consumption ($\Delta C_t^i < 0$). In the RIP model, this will never happen.

An example of this case is shown in figure 1. This graph plots the income paths of two individuals, where we continue to assume $\varepsilon_t^i \equiv 0$ for simplicity. Individual 1 experiences a faster average income growth rate in the first five periods than individual 2, but observes the same rise in income between periods five and six. If these income paths are generated by a RIP process (and individuals correctly perceives them as such), then both individuals will adjust their consumption

¹⁰It is true that if individuals could not separately observe z_t and ε_t in the RIP model but were solving a signal extraction problem instead, the history of shocks would also matter in the RIP model. However, the specific predictions implied by the HIP model described below would still not hold in such a model.

Figure 2: Distinguishing HIP from RIP (from Consumption Levels)




growth by exactly the same amount between periods five and six. Instead, if the truth is as in the HIP model, individual 1 will have formed a belief that his income growth rate is higher than that of individual 2, and was expecting his income to be closer to the trend line (shown by the dashed blue line). Therefore, even though his income increases, it is significantly below the trend ($\hat{\xi}_t < 0$), which causes him to revise down his beliefs about his true β^i , and consequently his consumption level from equation (15). Specifically, we have:

Prediction 1: The HIP model with Bayesian learning predicts that controlling for current income growth, consumption growth will be a *decreasing* function of average past income *growth rate*. Instead, the RIP model predicts no dependence on past income growth rate of this kind.

It is also possible to obtain a closed-form expression for the consumption (level) in the simplified version of the HIP model described above (Here we simply give an intuitive description of the information contained in the level of consumption, rather than going through the derivation). One can easily see that the *level* of consumption contains information about whether individuals perceive their income process as HIP or RIP. An example of this is shown in figure 2. This example is most easily explained when income shocks are permanent ($\rho = 1$), which we assume for the moment. As before, individuals realize different income growth rates up to period 3. Under the RIP model, both individuals' forecast of their future income is the same as their current income (shown with the horizontal dashed lines). In contrast with a HIP process, individual 1 will expect a higher income growth rate and therefore a much higher lifetime income than individual 2. Therefore, the first

Figure 3: Determining the Amount of Prior Knowledge in HIP



individual will have a higher consumption level than individual 2 at the same age, despite the fact that their current income levels are very similar. Therefore, we have:

Prediction 2: The HIP model predicts that controlling for the current level of income and past average income level, an individual's current consumption level will be an increasing function of his past income growth rate. This implication is independent of whether or not individuals know their true income growth rate.

Finally, it is also easy to see that the level of consumption is also informative about how much prior information individuals have about their own β^i within the HIP framework (question 3 raised in the introduction). To see this, consider the next figure (3) which is a slight variation of the previous one. Here both individuals are assumed to have observed the same path of income growth up to period 3 even though their true β^i are different. (This is possible since there are many stochastic shocks to the income process over time (coming from η_t), and the contribution of β^i to income is quite small). In this case, under the HIP model, if individuals have no private prior information about their own true β^i (which will be the case when $\lambda = 0$) then both individuals should have the same consumption level. The more prior information each individual has about his true β^i the higher will be the consumption of the first individual compared to the second. Therefore, an auxiliary model can capture this relationship by focusing on the following dynamic relationship:

Prediction 3: if $\lambda > 0$, then controlling for past income growth (as well as the current income level and past average consumption level) the consumption level of an individual

will be increasing in his *future* income growth as well. This is because in this case the individual has more information about his true β^i than is known to the econometrician and what is revealed by his past income growth.

These three examples illustrate how one can use the structural relationships that hold true exactly in a somewhat simplified version of the economic model of interest in order to come up with an auxiliary model. Indirect inference allows one to think in terms of these rich dynamic relationships instead of a set of moments (covariances, etc.). Below we are going to write a parsimonious auxiliary model that will capture these dynamic relationships to identify HIP from RIP and will also determine the degree of prior information (or equivalently, uncertainty) individuals face upon entering the labor market in the case of the HIP process.

3.1.1 A Parsimonious and Feasible Auxiliary Model

As shown above, the HIP model implies:

$$\Delta C_t = \Pi(\lambda, P_{t|t}, r, \rho, t, R, T) \times \left(Y_t^i - \left(\widehat{\beta}_{t|t-1}^i t + \widehat{z}_{t|t-1}^i \right) \right), \quad (17)$$

where $\Pi(\lambda, P_{t|t}, r, \rho, t, R, T) \equiv \Phi_{t,T}^r(A_t/X_t) + \Psi_{T-t}^{r,\rho}(B_t/X_t)$; the dependence of Π on λ and $P_{t|t}$, can be seen from the formulas for A_t and B_t . However, since $\widehat{\beta}_{t|t-1}^i$ and $\widehat{z}_{t|t-1}^i$ are unobserved by the econometrician (because they depend on all past income realizations as well as on each individuals' unobserved prior beliefs), this regression is not feasible as an auxiliary model. Moreover, we derived this relationship assuming a simplified HIP income process, quadratic utility, no borrowing constraints, and no retirement period, none of which is true in the life-cycle model we would like to estimate. Fortunately, as mentioned earlier, none of these issues represent a problem for the consistency of the estimates of the structural parameters that we are interested in.

We approximate the relationship in (17) with the following feasible regression:

$$c_t = a_0 + a_1 y_{t-1} + a_2 y_{t-2} + a_3 y_{t+1} + a_4 y_{t+2} + a_5 \overline{y}_{1,t-3} + a_6 \overline{y}_{t+3,T} \\ + a_7 \overline{\Delta y}_{1,t-3} + a_8 \overline{\Delta y}_{t+3,T} + a_9 c_{t-1} + a_{10} c_{t-2} + a_{11} c_{t+1} + a_{12} c_{t+2} + \epsilon_t$$

where c_t is the logarithm of consumption; y denotes the logarithm of labor income; $\overline{y}_{a,b}$ denotes the average of log income from time a to b ; and similarly $\overline{\Delta y}_{a,b}$ denotes the average growth rate of log income from time a to b . Notice that we use the logarithm of variables rather than the level; since the utility function is CRRA and income is log-Normal this seems to be a more natural specification. This regression captures the three predictions made by the HIP and RIP models discussed above by adding the past and future income growth rate as well as past and future income levels. To complete the auxiliary model we add a second equation with y_t as the dependent variable, and use all the regressors above involving income as left hand side variables (i.e., the nine

regressors excluding the lags and leads of consumption). Finally, we divide the population into two age groups—those between 25 and 38 years of age, and those between 39 and 55 years of age—and allow the coefficients of the auxiliary model to vary across the two groups.¹¹ For each age group, the auxiliary model has 22 regression coefficients (13 in the first equation and 9 in the second) and 3 elements in the covariance matrix of the residuals (one variance term for each equation and one covariance term between the two) for a total of 25 parameters. With two age groups, this yields a total of 50 reduced form parameters that determine the likelihood of the auxiliary model.

To implement the indirect inference estimator, we choose the values of the structural parameters so that the (approximate) likelihood of the observed data (as defined by the auxiliary model) is as large as possible. That is, given a set of structural parameters, we simulate data from the model, use this data to estimate the auxiliary model parameters, and evaluate the likelihood defined by the auxiliary model at these parameters. We then vary the structural parameters so as to maximize this likelihood. Viewed from another perspective, we are simply minimizing the difference between the (log) likelihood evaluated at two sets of auxiliary model parameters: the estimates in the observed data and the estimates in the simulated data (given a set of structural parameters). The advantage of this approach over other approaches to indirect inference (such as efficient method of moments or minimizing a quadratic form in the difference between the observed and simulated auxiliary model parameters) is that it does not require the estimation of an optimal weighting matrix. It is, however, less efficient asymptotically than the other two approaches, though this inefficiency is small when the auxiliary model is close to being correctly specified (and vanishes in the case of correct specification).

3.2 The Data

3.2.1 Constructing a Panel of (Imputed) Consumption

An important impediment to the previous efforts to use consumption data has been the lack of a sufficiently long panel on consumption expenditures. The Panel Study of Income Dynamics (PSID) has a long panel dimension but covers limited categories of consumption whereas the Consumer Expenditure Survey (CE) has detailed expenditures over a short period of time (four quarters). As a result, most previous work has either used food expenditures as a measure of non-durables consumption (available in PSID), or resorted to using repeated cross-sections from CE under additional assumptions.

In a recent paper Blundell, Pistaferri and Preston (2006b, BPP) develop a structural imputation

¹¹Although, the auxiliary model would correspond to the structural equation in (17) more closely if the coefficients were varying freely with age, this would increase the number of parameters in the auxiliary model substantially. Our experience is that the small sample performance of the estimator is better when the auxiliary model is more parsimonious, and therefore we opt for the specification here.

method which imputes consumption expenditures in PSID using information from CE. The basic approach involves estimating a demand system for food consumption as a function of nondurable expenditures, a wide set of demographic variables, and relative prices as well as the interaction of nondurable expenditures with all these variables. The key condition is that all the variables in the demand system must be available in the CE data set, and all but non-durable expenditures must be available in PSID. One then estimates this demand system from CE, and as long as the demand system is monotonic in nondurable expenditures, one can invert it to obtain a panel of imputed consumption in the PSID.

BPP implement this method to obtain imputed consumption in PSID for the period 1980 to 1992, and show that several statistics of the imputed consumption compare very well to their counterparts from CE. In this paper, we modify and extend the method proposed by these authors as follows. First, these authors include time dummies interacted with nondurable expenditures in the demand system to allow for the budget elasticity of food demand to change over time, which they find to be important for the accuracy of the imputation procedure. However, CE is not available on a continuous basis before 1980, whereas we would like to use the entire length of PSID from 1968 to 1992, making the use of time dummies impossible. To circumvent this problem, we replace the time dummies with other terms that are available throughout our sample period—specifically, the interaction of nondurable expenditures with food and fuel inflation rate. The inclusion of these inflation variables is motivated by the observation that the pattern of time dummies estimated by BPP after 1980 is similar to the behavior of these inflation variables during the same period.


A second important element in our imputation is the use of CE data before 1980. In particular, CE data are also available in 1972 and 1973, and in fact these cross-sections contain a much larger number of households than the waves after 1980.¹² The data in this earlier period are also superior in certain respects to those from subsequent waves: for example, as shown by Slesnick (1992), when one aggregates several sub-components of consumption expenditures in the CE, they come significantly closer to their counterparts in the National Income and Product Accounts than the CE waves after 1980.¹³ The use of this earlier data provides, in some sense, an anchor point for the procedure in the 1970's that improves the overall quality of imputation as we discuss further below.

Finally, instead of controlling for life-cycle changes in the demand structure using a polynomial in age (as done by BPP), we use a piecewise linear function of age with four segments, which provides more flexibility. This simple change improves the life-cycle profiles of mean consumption and the variance of consumption rather significantly. With these modifications, we obtain an

¹²The sample size is around 9500 units in 1972-73 surveys, but range from 4000-6000 units in the waves after 1980.

¹³For example, in 1973 total expenditures measured by the CE is 90 percent of personal consumption expenditures as measured by NIPA, whereas this fraction is consistently below 80 percent after 1980 and drops to as low as 75 percent in 1987. Similarly, consumer services in CE accounts for 93 percent of the same category in NIPA in 1973, but drops to only 66 percent in 1989.

Figure 4: Cross-sectional Variance of Log Consumption in CEX and Imputed PSID Data: 1968-1992.



imputed consumption measure that has a rather good fit to the statistics from CE as we discuss in a moment.

Since food and non-food consumption are jointly determined, some of the right hand side variables in the demand system are endogenous. In addition, nondurable expenditures are likely to suffer from measurement error (as is the case in most survey data sets), which necessitates an instrumental variables approach. We instrument log nondurable expenditures (as well as its interaction with demographics and prices) with the cohort-year-education specific average of the log of the husband's hourly wage and the cohort-year-education specific average of the log of the wife's hourly wage (as well as their interaction with the demographics and prices).

Table 1 reports the estimate of the demand system using the CE data. Several terms that include the log of nondurable expenditures are significant as well as several of the demographic and price variables. Most of the estimated coefficients have the expected sign. We invert this equation to obtain the imputed measure of household non-durable consumption expenditures.


Figure 4 plots the cross-sectional variance of log consumption over time. BPP used this figure as the main target to evaluate the satisfactoriness of their imputation. The line marked with squares shows the CE data whereas the line marked with circles shows the imputed consumption, which

Table 1: INSTRUMENTAL VARIABLES ESTIMATION OF DEMAND FOR FOOD IN THE CEX

Variable	Estimate	Variable	Estimate
$\ln(c)$	0.798*** (26.80)	$\ln(c) \times I\{11\% \leq \Delta \log p_{fuel}\}$	0.00386* (1.83)
$\ln(c) \times age \times I\{age < 37\}$	0.00036*** (3.38)	$\ln(c) \times (year - 1980)$	-0.00057 (-0.68)
$\ln(c) \times age \times I\{37 \leq age < 47\}$	0.00048*** (5.45)	One child	0.149 (1.16)
$\ln(c) \times age \times I\{47 \leq age < 56\}$	0.00042*** (5.75)	Two children	0.564*** (3.98)
$\ln(c) \times age \times I\{56 \leq age\}$	0.00037*** (6.08)	Three children+	1.203*** (8.23)
$\ln(c) \times$ High school dropout	-0.129*** (-7.57)	High school dropout	1.207*** (7.61)
$\ln(c) \times$ High school graduate	-0.043*** (-2.78)	High school graduate	0.417*** (2.90)
$\ln(c) \times$ One child	-0.014 (-1.01)	Northeast	0.0587*** (10.36)
$\ln(c) \times$ Two children	-0.055*** (-3.68)	Midwest	0.0293*** (5.23)
$\ln(c) \times$ Three children+	-0.123*** (-7.92)	South	-0.0031 (-0.63)
$\ln(c) \times I\{5\% \leq \Delta \log p_{food} < 8\%\}$	0.00096 (1.01)	Family size	0.0509*** (16.20)
$\ln(c) \times I\{8\% \leq \Delta \log p_{food} < 11\%\}$	0.00858*** (4.25)	$\ln p_{food}$	0.581** (2.28)
$\ln(c) \times I\{11\% \leq \Delta \log p_{food}\}$	-0.00091 (-0.39)	$\ln p_{fuel}$	-0.117 (-0.97)
$\ln(c) \times I\{5\% \leq \Delta \log p_{fuel} < 8\%\}$	0.00074 (0.66)	White	0.0824*** (11.38)
$\ln(c) \times I\{8\% \leq \Delta \log p_{fuel} < 11\%\}$	0.00091 (0.53)	Constant	-1.822*** (-2.65)
		Observations	21864

We pool the data from the 1972-73 waves of the CE with the 1980-92 waves. We instrument log food expenditures (and its interactions) with the cohort-education-year specific average of the log husband's and wife's hourly wage rates (and their interactions with age, education, and inflation dummies and a time trend). The t-statistics are contained in parentheses. The lowest value of Shea's partial R^2 for instrument relevance is 0.086, and the p-value of the F-test on the excluded instruments is smaller than 0.001 for all instruments.

Figure 5: Life-cycle Profile of Average Consumption in CEX and Imputed PSID Data




tracks the former fairly well, showing an overall rise in consumption inequality of 7-8 log points between 1980 and 1986, followed by a drop from 1986 to 1987 and not much change after that date. The dashed line shows that if one simply were to use food expenditures in PSID instead, the overall pattern would remain largely intact, but the movements would be *quantitatively* much more muted than in the data: the rise in consumption inequality would be understated by more than half by 1986 and by as much as two-thirds in 1991.

We also evaluate the quality of the imputation in two other dimensions that are important for the estimation exercise. First, figure 5 plots the average life-cycle profile of consumption implied by the CE data (line marked with squares) as well as the counterpart generated by imputed data (line marked with circles).¹⁴ The two graphs overlap remarkably well. The figure also plots the profile that is generated if we do not use the 1972-73 CE in the imputation procedure (dashed line): average consumption would rise by 51 percent between ages 25 and 45 instead of the 24 percent rise in the baseline imputation. Next, figure 6 plots the within-cohort variance of consumption over the lifecycle. Both in the CE and imputed PSID data, the variance rises almost in a linear fashion by about 10 log points between age 30 and 60.

It is also useful to provide some evidence on the quality of the imputation by testing its out-of-sample predictive ability at the household level. Specifically, we randomly split the CE sample used above into two subsamples (and make sure that each subsample contains exactly half of the

¹⁴The lifecycle profiles are obtained by controlling for cohort effects as described, for example, in Guvenen (2007b).

Figure 6: Life-cycle Profile of Consumption Variance in CEX and Imputed PSID Data



observations in each year). We use the first subsample to estimate the food demand system as above, which we then use to impute the non-durable consumption of the second subsample (control group). To eliminate sampling variation that results from the randomness of each subsample, we repeat this exercise 50 times. The discussions below refer to the average of these 50 replications. Comparing the actual non-durable expenditures of these households to that implied by the imputation is informative about the quality of the imputation. Figure 7 plots the actual consumption of the control group against the imputed one for each household (for the simulation with the median regression slope). The imputed consumption data forms a cloud that align very well along the 45-degree line. In fact a linear regression of imputed consumption on actual one yields an average slope coefficient of 0.996 and a constant term of 0.25. The average R^2 of the regression is 0.67, implying that the imputed consumption has a correlation of 0.81 with the actual consumption at household-level.¹⁵

The fact that the slope coefficient is almost equal to 1 is important: a slope above 1 would indicate that the imputation systematically overstates the variance of true consumption, which would in turn overstate the response of consumption to income shocks, thereby resulting in an overestimation of the size of income shocks. Clearly, the reverse problem would arise if the slope coefficient was below 1. Furthermore, when a quadratic term is added to the regression of imputed consumption on actual consumption, it almost always comes out as insignificant. This implies that

¹⁵ Across simulations, the slope coefficient in the regression range from 0.978 to 1.020, and the R^2 range from 0.644 to 0.691.




Figure 7: **Out of Sample Predictive Power of the Imputation Method in the CEX.** This plot is obtained by estimating the IV food demand system on a randomly chosen half of the CEX sample, and then imputing the consumption for the other half (control group). The figure plots the actual consumption of the control group versus their imputed consumption. The average regression slope is 0.996, the average constant is 0.24, and average R^2 is 0.67 over 50 repetitions.

the imputation procedure does not result in systematic under- or over-prediction at different points in the distribution, which would again be problematic. We next repeat the same exercise but now by only using the 1980 to 1992 waves of the CEX (figure 8). The results are very similar: the average slope is 1.036, the constant term is 0.12, and the R^2 is unchanged from before, at 0.67.

As a final, and quite strict test to detect systematic patterns in the imputation error, we regressed the difference between imputed and actual consumption for each individual (i.e., imputation error) on household characteristics including dummies for each age group, education dummies, family size, region dummies, number of children dummies and food and fuel prices. The median R^2 of this regression was 0.002 (and there was at most one variable that was significant at 5 percent level in each simulation) indicating no evidence of systematic imputation errors by demographic groups. Overall, we conclude that the imputation procedure works rather well and does not result in any systematic over- or under-prediction of actual consumption, which is reassuring for the estimation exercise we conduct in the next section.




Figure 8: **Out of Sample Predictive Power of the Imputation Method, CEX 1980-92.** This plot is obtained the same way as figure 7 but using only 1980-92 CEX data. The average regression slope is 1.03, the average constant is -0.12, and average R^2 is 0.67 over 50 repetitions.

Measure of Household Labor Income. In PSID, households report their total taxable income which includes labor income, transfers and financial income of all the members in the household. The measure of labor income we use subtracts financial income from this measure, and therefore, includes the labor income of the head and spouse as well as several categories of transfer income (unemployment benefits, social security income, pension income, worker’s compensation, welfare payments, child support, and financial help from relatives are the main components). PSID also reports estimated total taxes for all households until 1991. For 1992 and 1993 we use the TAXSIM software available from NBER to estimate taxes for each household. Since our income measure excludes asset income, for each year we regress this tax variable on the asset income and labor income of each household to back out the labor portion of the taxes paid in each year. We then subtract this estimated labor income tax from household income above to obtain the household after-tax labor income measure used in the analysis below.

Converting the Data to Per-adult Equivalent Units. We adjust both the imputed consumption and income measures for demographic differences across households since such differences have no counterpart in our model. This is accomplished by regressing each variable on family size, an education dummy, a race dummy, a number of children dummy, a region dummy, a dummy indicating whether the head is employed, the number of earners in the household, a dummy in-

Table 2: BASELINE PARAMETERIZATION FOR MONTE CARLO SIMULATION

Annual model		
Parameter		Value
δ	Time discount factor	0.945
r	Risk-free rate	0.04
ϕ	Relative risk aversion	2
R	Retirement age	41
$T - R$	Retirement duration	30

dicating residence in a large city, and a set of cohort dummies.¹⁶ We then use the residuals of these regressions—which are interpreted as consumption and income per-adult equivalent—in the analysis below.

4 Results

4.1 A Monte Carlo Exercise

In this section, we apply the proposed methodology to the estimation of the full HIP model with learning described above. To demonstrate the ability of this estimation method to uncover the true structural parameter vector in spite of (very) incomplete individual histories, substantial measurement error, potentially binding borrowing constraints, etc., we conduct a Monte Carlo study using 150 “observed” data sets drawn from the true data generating process.

The missing observations in the Monte Carlo study are chosen to be exactly the same as in the observed data. We include only individuals with at least five observations, for a total of about 2,200 individuals with an average of 12 observations on each. In the model, the interest rate is set to 4%, and the subjective time discount factor is 0.945. The number of years in the working life is set to 41, and the number of years in the retirement period is set to 30. The model incorporates a simplified Social Security system in which individuals receive, in each of the retirement periods, 30% of their income at age 65 (the last year of the working life). Individuals have isoelastic utility with coefficient of relative risk aversion equal to 2. The borrowing constraint is set close to the loosest possible value consistent with almost sure repayment of debts at the end of life, so that few individuals hit the borrowing constraint during their lifetimes. This setup is nearly identical to the one that will be used in the estimation using “real” data, so it is a good test of the performance of the proposed estimation methodology.

¹⁶Each cohort is defined by 5-year bands based on the birth year of each individual, such as those born between 1951 and 1955, 1956 and 1960, etc.

We next add classical measurement error to both consumption and income:

$$\begin{aligned} y_t^{i,*} &= y_t^i + u_t^{i,y}, \\ c_t^{i,*} &= c_t^i + \bar{u}^{i,c} + u_t^{i,c} \end{aligned}$$

where $y_t^{i,*}$ and $c_t^{i,*}$ are measured income and consumption of household i , respectively, and $u_t^{i,y}$ and $u_t^{i,c}$ are i.i.d random variables with zero mean over time. Notice that we also added a second term to consumption, $\bar{u}^{i,c}$, which is an individual fixed measurement error with zero mean in the cross-section. This fixed effect is needed for two reasons. First, and most importantly, recall that we regress both income and consumption on a full set of demographics to convert these variables into per-adult equivalent terms. However, one effect of this adjustment is that it introduces level differences between consumption and income, the magnitude of which varies by household. This fixed effect captures such differences.¹⁷ Second, the model described above abstracts from initial wealth differences across households, which clearly exist in the data. These differences in wealth would also drive a household-specific wedge between the levels of income and consumption. The fixed effect is also a simple way to capture these differences in initial wealth levels. However, because all households in the simulated data have the same demographics and zero initial wealth, this fixed effect is redundant in the Monte Carlo analysis. Therefore, we set it to zero until we get to the estimation with real data below.

Incomplete histories are handled by “filling in” missing values in a reasonable way: in particular, for each individual we calculate the lifetime average of either log consumption or log income using available observations. If a consumption or income observation is missing in a given year, we simply replace the missing data with this average. We construct the past growth rate for age t in the auxiliary model by taking the difference between the latest valid observation before t and the first valid observation for the individual in data set and dividing this difference by the number of years between the two ages. The future growth rate at a given age is constructed analogously. If either variable cannot be constructed for a given age we use the average growth rate of that variable in the population instead. We use exactly the same procedure in both the simulated and observed data. The (approximate) likelihood, however, includes contributions only from those time periods in which the left-hand side variables are observed (i.e., not missing). Below, we consider alternative methods for filling in missing data and check the sensitivity of the results to the method used.

The results are contained in Table 3. The “true values” for the parameters are set to the estimates obtained in the next section using PSID labor income data (column 1 in table 4). The

¹⁷To see this, consider two households with the same income and consumption, but suppose that the first household has more children than the second, and both households have the same number of earners. Converting the variables to per-adult equivalent units will result in the first household having a lower consumption than the second one despite having the same income (since children consume but typically do not earn income). A similar issue arises between households with different number of earners given a certain level of total income and consumption. Since we do not explicitly account for such demographic differences in our model—which would complicate the analysis tremendously—we account for such differences in levels using the fixed effect as modeled here.

Table 3: ESTIMATING THE FULL CONSUMPTION-SAVINGS MODEL: MONTE CARLO RESULTS

	True value	Mean estimate	Std. dev.
ρ	0.670	0.669	0.034
σ_η	0.190	0.191	0.009
σ_ε	0.150	0.150	0.015
$\sigma_\beta(\times 100)$	2.650	2.652	0.121
σ_α	0.490	0.493	0.031
λ	0.500	0.504	0.054
σ_{u^y}	0.200	0.199	0.003
σ_{u^c}	0.200	0.199	0.008

Note: Statistics are based on 150 replications

initial values of the parameters are set randomly to $\pm 20\%$ of the true values. Each Monte Carlo run takes about 30 minutes on a state-of-the-art workstation. Clearly, the estimation method works well: bias is virtually absent and standard deviations are small. Although it is difficult, if not impossible, to prove identification in this setup, the results suggest strongly that local identification near the true parameter vector does indeed hold. These results are very encouraging and suggest strongly that the proposed methodology is a feasible and practical method for estimating structural consumption-saving models with missing data and multiple sources of heterogeneity.

4.2 Results using PSID data

We now estimate the lifecycle model using the PSID household after-tax labor income data and the imputed consumption data. In addition to the parameters above, we now also estimate the standard deviation of the consumption fixed measurement error, $\bar{u}^{i,c}$.

Table 4 reports the results. In the first column, we estimate the parameters of the income process using income data only. The estimated persistence is 0.64—far away from a unit root—with an innovation standard deviation of 19 percent per year. The transitory shock (which is a combination of genuine shocks and i.i.d measurement error) has a standard deviation of 14.6 percent. One of the main parameters of interest, σ_β , is estimated to be 2.65 percent, which is substantial. The dispersion in initial income levels, σ_α , is almost 50 percent, and has a negative correlation with individual income growth rate (-0.475). These results are consistent with the earlier studies in the literature that use representative panels on income and do not impose a priori restrictions on σ_β (Baker 1997, Haider 2001, Guvenen 2007b).¹⁸

¹⁸Notice, however, that we are not yet allowing for time-effects in the innovation variances which is common in this literature. So these results are not exactly comparable to these earlier studies at this point. This extension will be introduced in the next version of this paper.

Table 4: ESTIMATING THE FULL CONSUMPTION-SAVINGS MODEL USING REAL DATA

Data:	Y only	Y & C	Y &C	Y &C
		Baseline	$\sigma_\beta \equiv 0$ imposed	Alt. filling
ρ	0.636	0.789	0.972	0.773
σ_η	0.192	0.171	0.140	0.179
σ_ε	0.146	0.053	0.085	0.033
$\sigma_\beta (\times 100)$	2.645	1.877	—	1.78
σ_α	0.496	0.329	0.369	0.279
$\sigma_{\alpha\beta}$	-0.475	-0.303	—	-0.506
λ	—	0.768	—	0.795
σ_{u^y}	—	0.153	0.138	0.150
σ_{u^c}	—	0.341	0.357	0.340
$\sigma_{\bar{u}^c}$	—	0.402	0.385	0.391

We next estimate the full model (second column). The estimated persistence is now higher, at 0.79, but still well below a unit root with an annual standard deviation of 0.17. Therefore, the joint dynamics of consumption and income data do not lend support to permanent shocks as a reasonable representation of the typical income shock. With consumption data, in principle, we can tell apart transitory shocks from measurement error in income, since consumption should respond to the former but not to the latter. In practice, however, because the response of consumption to transitory shocks is proportional to its annuitized value—which is small—this response is rather weak and identification is a problem empirically. In this framework, however, borrowing constraints are binding for a significant fraction of households—no less than 20 percent of households younger than 35 years of age. As a result, these households’ consumption move one for one with transitory shocks allowing us to distinguish these shocks from pure measurement error. We estimate the standard deviation of transitory shocks to be about 5.3 percent and the standard deviation of measurement error in income to be about 15 percent.

The dispersion in income growth rates is now estimated to be 1.87 percent which, although smaller than in column 1, is still substantial. To understand the economic significance of this estimate, note that assuming a 1 percent average income growth rate per year, an individual who is one standard deviation above the mean will earn 2.1 times the median income and 4.4 times the income of an individual who is one standard deviation below the mean. Of course, not all this heterogeneity represents uncertainty on the part of the individuals, since each individual has some prior information about his true β^i by the time he enters the labor market. The parameter λ measures this prior information and is estimated to be 0.768. The interpretation is that the standard deviation of the prior uncertainty faced by individuals is $\sigma_\beta \times (1 - \lambda)^{1/2} \approx 1.87 \times 0.48 = 0.90$. Now consider an individual who enters the labor market with these prior beliefs and forms optimal

forecasts of his income all the way until retirement age. The upper bound of one-standard deviation confidence band for income at retirement age will be 105 percent larger than the lower bound. By comparison, the upper bound of the one-standard deviation confidence band resulting from persistent shocks is 72 percent larger than the lower bound. This calculation shows that the perceived income risk due to the uncertainty about income growth is substantial and is, in fact, the major component of lifetime income risk. Finally, the classical measurement error in consumption has a standard deviation of 35 percent and includes the noise introduced by the imputation method. Furthermore, the fixed effect in measured consumption (that results from the conversion to per-adult equivalent terms) has a standard deviation of 40 percent.¹⁹

We next examine the consequences of imposing an a priori restriction on the heterogeneity in income growth rates by setting $\sigma_\beta \equiv 0$ as done in the literature that estimates versions of the RIP process. As seen in column 3, the estimated persistence is now 0.97 with an innovation standard deviation of 0.14. Perhaps, the substantially higher persistence found in this case should not come as a big surprise. This is because Guvenen (2007b) shows that ignoring profile heterogeneity, when in fact it is present, leads to an upward bias in estimated persistence when only income data is used. The results we find here show that this is true even when consumption data is used in addition to income data.

Before closing this section it is useful to examine the robustness of these results to the method chosen for filling in missing data. This could be potentially important because more than two-thirds of the data in our sample is missing—and therefore filled in—compared to a fully balanced panel with the same number of individuals. In the last column we re-estimate the full model by using an alternative procedure to “fill in” missing observations.²⁰ As can be seen here, the results are largely unchanged compared to the baseline case in column 2, even though the two procedures are quite different. Therefore, this suggests that there is a good deal of flexibility in the method one uses to fill in missing data, which is reassuring.

5 Conclusions

The joint dynamics of consumption and labor income contains rich information that allows a sharper distinction between the RIP and HIP models. Monte Carlo results suggest that the indirect in-

¹⁹We found this component to be important for the overall estimation—failing to include this term results in implausible estimates for many parameter values, as the minimizer struggles to make sense of the fixed level differences in the data that has no counterpart in the model.

²⁰Basically, for each age an individual has a real data point, we find the percentile ranking of this observation in the entire distribution of income for that age in our sample. We then take the average of the percentile rankings for this individual over all the ages when he has a valid observation. Then for each missing observation of this individual, we simply impute the income level corresponding to his average percentile ranking given the income distribution in our sample for that age.

ference method works very well, even in the presence of frequently binding borrowing constraints, missing observations, retirement income, and so on, that make the auxiliary model a poor approximation to the structural relationships that need to hold in the model. On a more substantive level, we find that (i) income shocks have modest persistence, much less than a unit root, (ii) income growth rates display significant cross-sectional heterogeneity, (iii) individuals have much better information about their own income growth rates than what can be predicted by some observable variables available to the econometrician, and (iv) finally, despite significant prior information, there is also large prior uncertainty that affects consumption behavior throughout the lifecycle.

The results reported so far should be viewed as a progress report on this ongoing project. We plan to extend and enrich this basic framework in several directions. The next immediate step to introduce time effects in variances and non-separable leisure into the utility function. We also plan to model partial insurance over and above what can be achieved by self-insurance. By imposing sufficient theoretical structure on the form of this partial insurance, it seems possible to disentangle it from the other sources of uncertainty and insurance available in the model. Finally, we also aim to conduct formal statistical tests to assess the extent to which the RIP model can be rejected against the alternative of HIP using consumption and income data jointly.

References

- [1] Abowd, John, and David Card (1989): “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, Vol. 57, pp.411-445.
- [2] Baker, Michael (1997): “Growth-rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings,” *Journal of Labor Economics*, Vol. 15, pp. 338-375.
- [3] Becker, Gary S. (1965): “Human Capital : A Theoretical and Empirical Analysis, with Special Reference to Education,” reprinted 1994, University of Chicago Press.
- [4] Ben-Porath, Yoram (1967): “The Production of Human Capital and the Life Cycle of Earnings,” *Journal of Political Economy*, Vol. 75, pp 352-365.
- [5] Blundell, Richard, and Ian Preston (1998): “Consumption inequality and income uncertainty,” *Quarterly Journal of Economics*, Vol. 113, pp. 603-640.
- [6] Blundell, Richard, Luigi Pistaferri, and Bruce Preston (2006a): “Consumption Inequality and Partial Insurance,” Stanford University Working Paper.
- [7] Blundell, Richard, Luigi Pistaferri, and Bruce Preston (2006b): “Imputing Consumption in the PSID Using Food Demand Estimates in the CE,” Stanford University Working Paper.
- [8] Cunha, Flavio, James Heckman, and Salvador Navarro (2005): “Separating uncertainty from heterogeneity in life cycle earnings,” *Oxford Economic Papers*, forthcoming.

- [9] Deaton, Angus, and Christina Paxson (1994): “Intertemporal Choice and Inequality,” *Journal of Political Economy*, Vol. 102, pp. 437-467.
- [10] Gourinchas, Pierre-Olivier, and Jonathan Parker (2002): “Consumption Over the Life Cycle,” *Econometrica*, Vol 70, pp. 47-89.
- [11] Guvenen, Fatih (2007a): “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?” *American Economic Review*, Vol. 97, No. 3, pp. 687-712.
- [12] Guvenen, Fatih (2007b): “An Empirical Investigation of Labor Income Processes,” National Bureau of Economic Research Working Paper No 13394.
- [13] Guvenen, Fatih and Burhanettin Kuruscu (2006): “A Quantitative Analysis of the Evolution of the U.S. Wage Distribution: 1970—2000,” Working Paper, University of Texas at Austin.
- [14] Haider, Steve J. (2001): “Earnings Instability and Earnings Inequality of Males in the United States: 1967-1991,” *Journal of Labor Economics*, Vol. 19, pp. 799-836.
- [15] Hall, Robert E., and Frederik S. Mishkin (1982): “The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households,” *Econometrica*, Vol 50, No 2, pp. 461-81.
- [16] Hause, John C. (1980): “The Fine Structure of Earnings and the On-the-Job Training Hypothesis,” *Econometrica*, Vol. 48, pp. 1013-29.
- [17] Heathcote, Jonathan, Kjetil Storesletten, Gianluca Violante (2006): “The Macroeconomic Implications of Rising Wage Inequality in the United States,” Working Paper, NYU.
- [18] Huggett, Mark, Gustavo Ventura, and Amir Yaron (2006): “Sources of Lifetime Inequality,” Working Paper, Wharton School, University of Pennsylvania.
- [19] Lillard, Lee A., and Yoram A Weiss (1979): “Components of Variation in Panel Earnings Data: American Scientists, 1960-70,” *Econometrica*, Vol. 47, pp. 437-454.
- [20] MaCurdy, Thomas (1982): “The Use of Time-Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis,” *Journal of Econometrics*, Vol. 18, pp. 83-114.
- [21] Meghir Costas and Luigi Pistaferri (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, Vol. 72, pp. 1-32.
- [22] Moffitt, Robert A., and Peter Gottschalk (1994): “Trends in the Autocovariance Structure of Earnings in the US, 1969-1987, Working Paper, Brown University.
- [23] Pischke, Jörn-Steffen (1995): “Individual Income, Incomplete Information, and Aggregate Consumption,” *Econometrica*, Vol. 63, pp. 805-840.

- [24] Slesnick, Daniel (1992): "Aggregate Consumption and Saving in the Postwar United States," *Review of Economics and Statistics*, Vol 74, No 4, pp. 585-97.
- [25] Topel, Robert H. (1990): "Specific Capital, Mobility and Wages: Wages rise with Job Seniority," NBER Working Paper No. 3294.